TENTH EDITION

Business Data Networks and Security

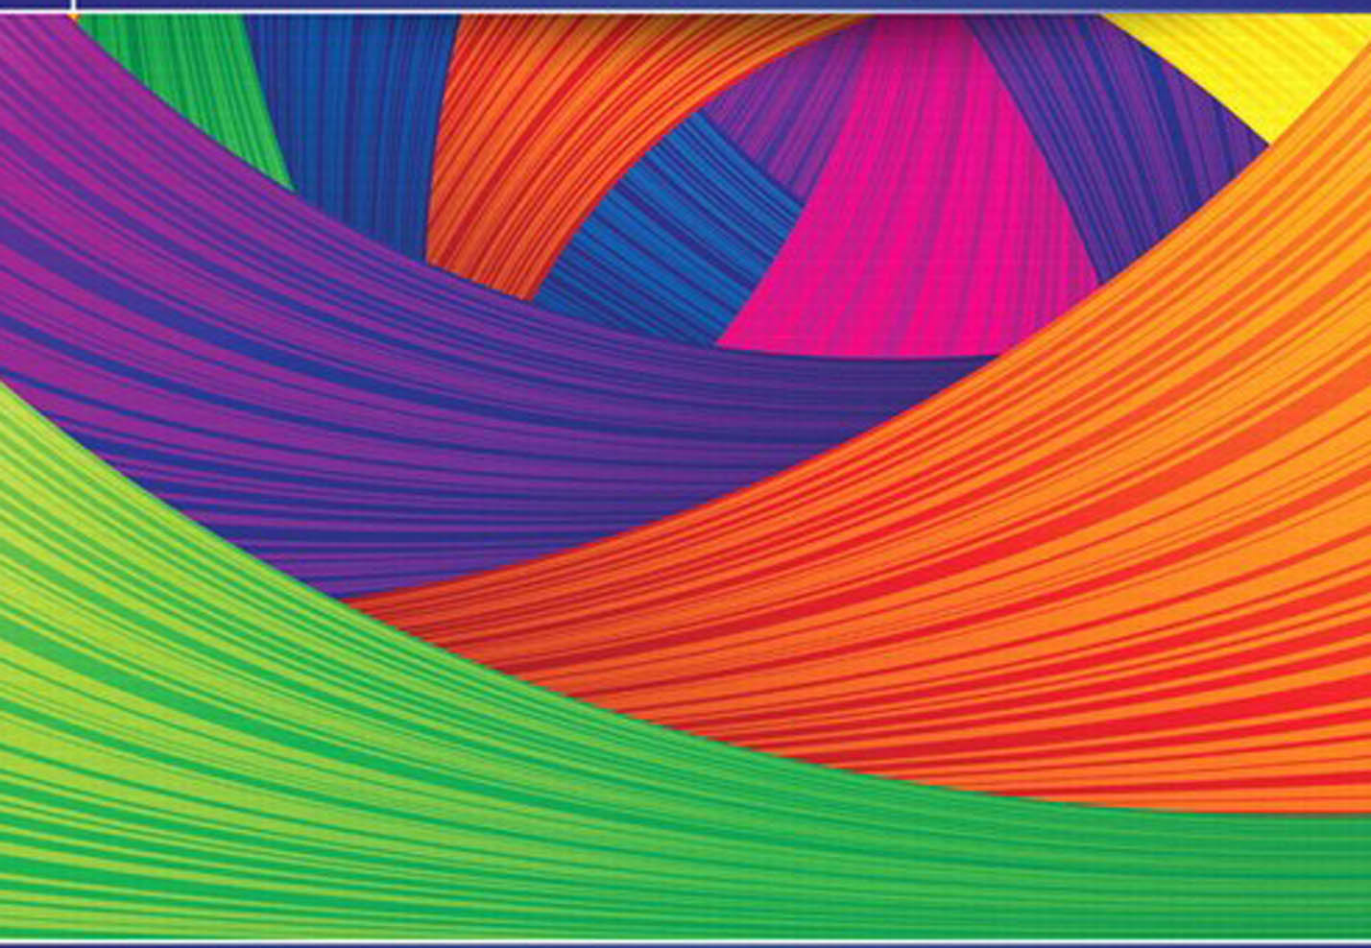

Raymond R. Panko • Julia L. Panko

Business Data Networks and Security

This page intentionally left blank

Tenth Edition

Business Data Networks and Security

Raymond R. Panko
University of Hawai`i at Mānoa

Julia L. Panko
Weber State University

PEARSON

Boston Columbus Hoboken Indianapolis New York San Francisco
Amsterdam Cape Town Dubai London Madrid Milan Munich Paris Montreal Toronto
Delhi Mexico City Sao Paulo Sydney Hong Kong Seoul Singapore Taipei Tokyo

Editor-in-Chief: Stephanie Wall
Director of Marketing: Maggie Moylan
Executive Marketing Manager: Anne Fahlgren
Project Manager: Tom Benfatti
Acquisitions Editor: Nicole Sam
Program Manager: Denise Vaughn
Program Manager Team Lead: Ashley Santora
Project Manager Team Lead: Judy Leale

Cover Designer: Jon Boylan, Lumina
Cover Image: margouillat/fotolia.com
Full Service Project Management: Allan Rayer
Composition: Integra Software Solutions
Printer/Binder: Edwards Brothers Malloy
Cover Printer: Lehigh-Phoenix Color, Hagerstown
Text Font: 10/12 Palatino LT Std

Library of Congress Cataloging-in-Publication Data On File

10 9 8 7 6 5 4 3 2

PEARSON

ISBN 10: 0-13-354401-X
ISBN 13: 978-0-13-354401-5

To Sal Aurigemma. A great partner in crime in research and teaching.

BRIEF CONTENTS

Online Modules

(available at www.pearsonhighered.com/panko)

CONTENTS

Online Modules

(available at www.pearsonhighered.com/panko)

Module A MORE ON TCP

Module B MORE ON MODULATION

Module C MORE ON TELECOMMUNICATIONS

This page intentionally left blank

PREFACE FOR STUDENTS

Networking and security are the most exciting careers in information technology. Heck, they are the most exciting careers in the world. Professionals in these fields do not spend their careers just doing the same thing over and over again. Their work is constantly evolving, and personal growth is guaranteed.

HOW TO STUDY NETWORKING

Networking and Security Are Different

Some students find networking and security difficult. The problem seems to be that they require a different learning approach than programming and database management. In programming and database, you learn a little, apply it, learn a little more, apply it, shampoo, rinse, repeat. If there is something you don't know, there is probably another way to do it. (Except on exams and homework, of course.)

In networking, you need to know everything to do anything, and it is what you don't know that hurts you. For example, suppose that you want to connect a server to an Ethernet switch. This sounds simple enough. However, should you choose copper wire or optical fiber? If copper wire, what grade of copper wire? If fiber, which OM standard should you choose? Or should you connect the server wirelessly? In your choice, you must include speed, distance, delay, reliability, and cost. Especially cost. Budgets are eternally tight, and networking people never say "cost doesn't matter."

Security is different again. In security, you are not just dealing with design issues and the reliability of technology. You are dealing with human opponents that are engaged with you in a perpetual arms race of protections and new attack methods to get beyond those protections. It is a lot like playing a video game at a high level, but with real-world consequences.

Will employers expect you to know everything when you apply for a job? Of course not. However, they will expect you to know a *lot*. They will sit you down and ask you how to connect a server to an Ethernet switch or something else that requires you to be able to integrate what you have learned. In fact, they will do this for the material in most courses you have taken to get an understanding of how serious you are about work.

You will certainly get questions that require you to troubleshoot a problem. Troubleshooting is hard, and most people intuitively do it wrong. This book will give you a methodology for doing it right and plenty of practice in applying it.

Employers will expect applicants to be up in the field. For Wi-Fi, they may ask you about security, and they don't expect you to stop at 802.11i. Mentioning Ethernet busses and hubs in a design may end the interview. Employers expect applicants to have some knowledge of IPv6 and cloud computing. They will be interested if you know even a little about SDN.

Learning with this Book

Organization of the Book We have tried to write this book to help you learn the material. Most basically, we present the material in short sections with Test Your Understanding (TYU) questions immediately after each section, to help you know if you have understood the section.

Pay special attention to keyterms that are boldfaced. These are the core concepts in the field. And yes, there are a lot of them. Important or frequently misunderstood concepts are broken out like this for special attention:

A rogue access point is an unauthorized access point set up within a firm by an employee or department.

Figures cover almost all important concepts in the book. There are special study figures that summarize the flow and key points in most sections that are not amenable to illustrations. The PowerPoint presentations are based on these figures. For complex illustrations, the PowerPoint presentations have builds, presenting only part of the figure at each step.

If you see a term that you learned previously but have forgotten, go to the Glossary. In Glossary entries, some page numbers are boldfaced. These are the pages on which the term was defined or characterized. Some terms are introduced more than once and may have two or more page numbers boldfaced.

Studying for Exams Exams are the least popular elements in any course. And yes, you will have dreams about waking up late for an exam for several years after you graduate. However, there are things you can do to make your life easier.

First, study the material. Read a section. Do the TYU questions. In fact, download the homework file, which has all the questions. Put your answers into the file. The multiple choice questions in the test bank are taken from the material in the TYU questions and thought questions. A good idea is to read the material over before exams instead of just relying on your initial answers, which might not have been exactly perfect, having been based on your first reading.

Late in your study, describe the figures as if you were giving a lecture. If there is something you do not understand, note it and follow up. Take notes on your problems and insights.

At each step, ask yourself why each question and answer is important. This will give you insights and will solidify the material in your memory.

Upper-Division Learning Initial college education focuses on learning isolated facts. Networking and security, like other advanced courses, requires something more. First, it requires the ability to compare and contrast concepts you have learned. In networking and security, there are alternative ways to do almost everything. Understanding individual alternatives is not enough. To select the best alternative, you must understand trade-offs between them. You must also see them in the broader context of the chapter. For 802.11 Wi-Fi, 802.11i provides a lot of protection; but there

are other things you must also do to be secure. Life is about trade-offs. Your studying must reflect that.

Another pain point is learning multi-step procedures. It is important to learn the overall flow, understand how each step relates to the flow, understand each step, and do this all over again until you have both the flow and the details. Processes are difficult to learn because you do not have a framework clearly in mind for fitting individual facts into the bigger picture. In learning processes, it takes several cycles of studying at multiple levels to get both the overall flow and the individual steps.

ABOUT THE AUTHORS

Ray Panko is a professor of IT management and a Shidler Fellow at the University of Hawai'i's Shidler College of Business. His main courses are networking and security. Before coming to the university, he was a project manager at Stanford Research Institute (now SRI International), where he worked for Doug Englebart, the inventor of the mouse and creator of the first operational hypertext system. He received his B.S. in physics and his M.B.A. from Seattle University. He received his doctorate from Stanford University, where his dissertation was conducted under contract to the Office of the President of the United States. He has been awarded the Shidler College of Business's Dennis Ching award as the outstanding teacher among senior faculty. His e-mail is Ray@Panko.com.

Julia Panko is an assistant professor on the faculty at Weber State University. She received her doctorate from the University of California, Santa Barbara. Her research interests include the twentieth- and twenty-first-century novel, the history and theory of information technology, and the digital humanities. Her dissertation focused on the relationship between information culture and modern and contemporary novels.

Chapter 1

Welcome to the Cloud

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Describe basic networking, including why networks are drawn as clouds, hosts, addresses, the Internet, Internet service providers, transmission speed, and service level agreements.
- Explain how the Internet works, how Netflix uses Amazon Web Services IaaS (Infrastructure as a Service) with virtual machines, and a Google SaaS (Software as a Service).
- Describe messages, fragmentation, multiplexing, and frames versus packets.
- Describe how single point-to-point, wireless, switched, and hybrid wireless-switched networks operate—especially how switches forward incoming frames.
- Describe how internets and routers make it possible for hosts on different networks to work together.
- List the five standards layers commonly encountered in networking, describe what each layer does, describe concepts and terms in each layer, identify at which layer a given process is operating, and identify which standards agencies and standards architecture are relevant to that process.

BOX 1

By the Numbers

The Internet is enormous, growing, and changing.

- By 2003, there were already more devices connected to the Internet (computers, phones, etc.) than there were human users.[1]
- In 2010, 21% of the world's population used the Internet. In 2013, it was 39%.[2]
- In 2012, online video viewing overtook DVD and Blu-Ray viewing.[3]
- From 2011 to 2016, global IP traffic will triple, and the number of connected devices will nearly double.[4]
- In 2016, Cisco expects the Internet to carry one zettabyte of data.[5] A zettabyte is 1,000,000,000,000,000,000 (one sextillion) bytes.
- By 2020, there will be 50 billion devices connected to the Internet—ten times the number of human users. The great majority of these will be devices talking to other devices, without human involvement.[6]

NETFLIX DIVES INTO THE AMAZON[7]

Figure 1-1 shows that the Internet is often depicted as a cloud. This symbolizes that just as you cannot see inside a cloud, users should be oblivious to what happens inside the Internet. To them, the Internet simply works, like the electrical, water, and telephone systems.

In this course, as you might suspect, you will not be spared the burden of understanding the internals of the Internet and other networks. This knowledge will prepare you to help your employer use networks effectively. Along the way, you will learn a good deal about security, too. Networking is a vast superhighway with great potential for benefits. However, it has some rough neighborhoods.

---

[1] Suzanne Choney, "US Has More Internet-Connected Gadgets Than People," *nbcnews.com*, January 2, 2003. http://www.nbcnews.com/technology/us-has-more-internet-connected-gadgets-people-1C7782791.

[2] Geneva, "Key ICT Indicators for Developed and Developing Countries and the World (Totals and Penetration Rates)," *International Telecommunications Unions (ITU)*, February 27, 2013.

[3] Jared Newman, "Online Video Expected to Overtake DVD, Blu-ray Viewing this Year," *Techhive*, May 27, 2012. http://www.techhive.com/article/252650/online_video_expected_to_overtake_dvd_blu_ray_viewing_this_year.html.

[4] Larry Hettick, "Cisco: Networked Devices Will Outnumber People 3 to 1 in 2016," *Network World*, June 1, 2012. http://www.networkworld.com/newsletters/converg/2012/060412convergence1.html

[5] Grant Gross, "Cisco: Global 'Net Traffic to Surpass 1 Zettabyte by 2016, Cisco Says," *Network World*, May 31, 2012. http://www.pcworld.com/article/256522/cisco_global_net_traffic_to_surpass_1_zettabyte_in_2016.html

[6] Ericsson, "CEO to Shareholders: 50 Billion Connections 2020," press release, April 2010.

[7] Sources for this section include the following. Brandon Butler, "Three Lessons from Netflix on How to Live in a Cloud," *NetworkWorld*, October 9, 2013. http://www.networkworld.com/news/2013/100913-netflix-cloud-274647.html. Matt Petronzio, "Meet the Man Who Keeps Netflix Afloat in the Cloud," *mashable.com*, May 13, 2013. http://mashable.com/2013/05/13/netflix-dream-job/. Kevin Purdy, "How Netflix is Revolutionizing Cloud Computing Just So You Can Watch 'Teen Mom' on Your Phone," *www.itworld.com*, May 10, 2013. http://www.itworld.com/cloud-computing/355844/netflix-revolutionizing-computer-just-serve-you-movies. Ashlee Vance, "Netflix, Reed Hastings Survive Missteps to Join Silicon Valley's Elite," *Business Week*, May 9, 2013. http://www.businessweek.com/articles/2013-05-09/netflix-reed-hastings-survive-missteps-to-join-silicon-valleys-elite.

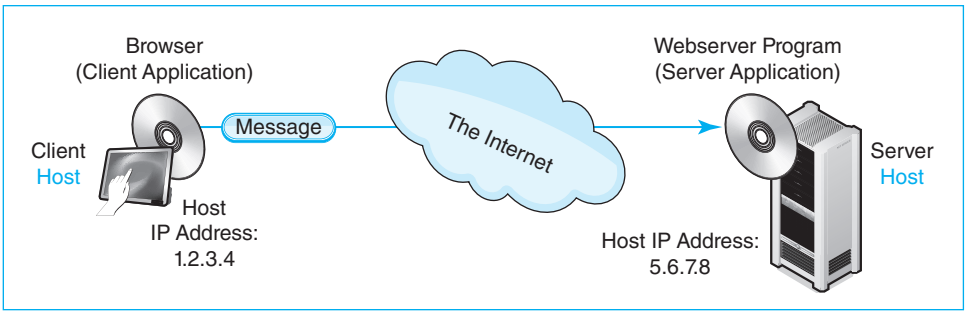

FIGURE 1-1 Internet Communication

Test Your Understanding

1. a) Why is the Internet usually depicted as a cloud? b) What is the significance of this depiction for users?

Hosts, Messages, and Addresses

Hosts Figure 1-1 introduces some basic networking terms. First, any computer attached to a network is a **host**. Hosts include large servers that work with hundreds of users simultaneously. Hosts also include desktop PCs, laptops, tablets, smartphones, smart glasses, and smart watches. In the future, hosts will include interactive walls, tables, and appliances that will turn your entire home into an immersive interactive environment. In a trend called the **Internet of things**, even coffee makers, toasters, medical implants, and many other small and large devices around us will be hosts that communicate through networks to work better. In fact, machine-to-machine communication will eventually dominate traffic on the Internet. The term *host* is not an obvious name for computers that attach to networks, but it is the common name for them in networking.

Any computer attached to a network is a host.

Messages and Addresses Figure 1-1 shows that application programs on different hosts communicate by sending messages to one another. Messages require addresses. For example if you want to send the first author a message, you would send it to his e-mail address, Ray@Panko.com. Hosts also need addresses. On the Internet, these are **Internet Protocol addresses** or **IP addresses**. In Figure 1-1, the IP addresses are 1.2.3.4 for the source host and 5.6.7.8 for the destination host.

Dotted Decimal Notation (DDN) When an IP address is expressed as four numbers separated by dots (periods), this is called **dotted decimal notation (DDN)**. In reality, IP addresses are 32-bit strings of 1s and 0s. Computers have no problem working with long bit strings. Human memory and writing, however, need a crutch to deal with long bit strings. Dotted decimal notation is precisely that—a crutch for inferior biological entities like ourselves. Computers do not use DDN.

32 IP address bits divided into four 8-bit segments	00000001	00000010	00000011	00000100
Segment converted to decimal	1	2	3	4
IP address in dotted decimal notation (DDN)	1.2.3.4			

FIGURE 1-2 Dotted Decimal Notation

Figure 1-2 shows how to convert a 32-bit IP address into dotted decimal notation.

- First, divide the 32 bits into four 8-bit segments.
- Second, treat each segment as a binary number and convert this binary number into a decimal number. For example, the first segment, 00000001 in binary, is 1 in decimal.
- Third, combine the four decimal field values, separating them by dots. This gives 1.2.3.4.

How do you convert a binary number into a decimal number? The fastest way is to go to an Internet search engine and find a binary-to-digital converter. You then type each 8-bit binary segment's bits into the indicated binary box and hit the convert button. The decimal value appears in the decimal box.

We have been looking at a 32-bit IP address. However, this is not the only type of IP address. It is an **IP Version 4 (IPv4)** address. IPv4 is the dominant IP protocol on the Internet today. However, we are beginning to see significant use of **IP Version 6 (IPv6)**. As we will see in Chapter 8, IPv6 addresses are 128 bits long and are represented for human consumption in a very different way.

Test Your Understanding

2. a) What is the term we use in networking for any computer attached to a network? b) Is your smartphone a host when you use it to surf the 'Web? c) Are you as a person a host when you use a network? d) How do application programs on different hosts communicate?
3. a) What kind of addresses do hosts have on the Internet? b) What kind of address is 128.171.17.13? c) What name do we use for the format 128.171.17.13? d) Who uses this format—humans or computers? e) Convert the following 32-bit binary IP address into DDN (spaces are added for easier reading): 10000000 10101011 00010001 00001101. (Check Figure: 10000000 = 128) f) Convert 5.6.7.138 into a 32-bit IP address. (Check Figure: 5 = 00000101) Show a space between each 8-bit segment. g) What type of IP address is 32 bits long? h) What other type of IP address exists, and how long are its addresses?

The Internet

Figure 1-3 illustrates that the global Internet is not a single network. Instead, the **Internet** is a collection of thousands of single networks and smaller internets. All of these single networks and smaller internets interconnect to form a single transmission system that in

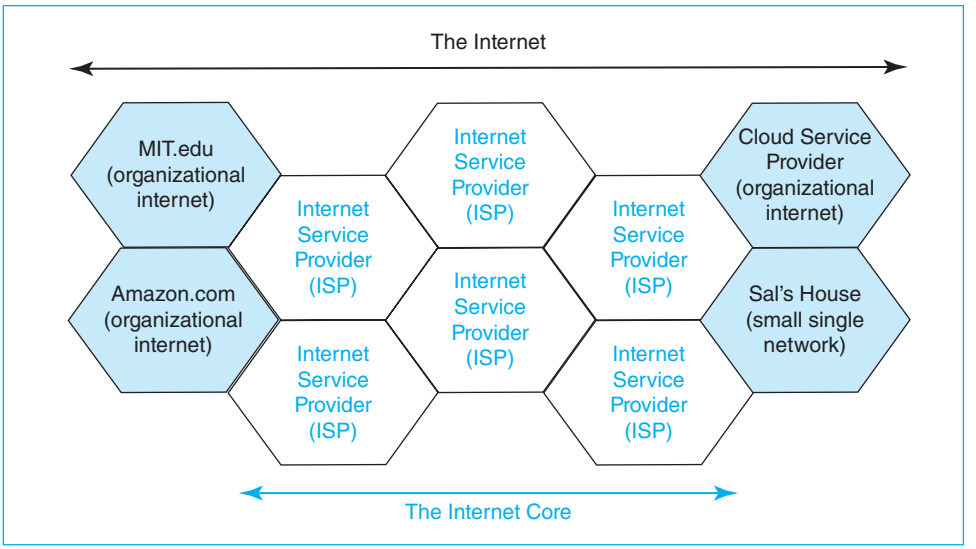

FIGURE 1-3 The Internet's Networks and Smaller Internets

principle allows any Internet host to reach any other.[8] Some of these single networks and smaller internets are owned by organizations such as Amazon.com or MIT. Smaller networks are owned by families and even individuals. In addition, some internets link these smaller networks and smaller internets together. We call these linking internets **Internet service providers (ISPs)**. ISPs collectively form the **core** of the Internet, which is also called the Internet's backbone.[9] To use the Internet, a customer must connect to an ISP.

The Internet is a collection of single networks and smaller internets. All of these networks and smaller internets interconnect to form a single transmission system.

At this point, we need to break the narrative to mention in two pieces of terminology we will use in this book.

- First, saying "single networks and internets" is cumbersome. We use the term *network* for both.
- Second, in this book, we spell internet in lowercase for internets in general and internets that are not the global Internet. We capitalize the global Internet.

Who owns the Internet? The surprising answer is, "Nobody." The ISPs and other organizations own their pieces of the Internet. Who controls the Internet? Again, nobody does. Although the **Internet Engineering Task Force (IETF)** creates standards,

[8] The original term for *internet* was *catanet*. When things are connected together in computer science, they are said to be concatenated. Fortunately, "catanet" never caught on, saving the Internet from a flood of bad feline jokes.

[9] For simplicity, the figure shows ISPs as if they served nonoverlapping geographic regions. Actually, ISPs often overlap geographically. National and international ISPs may connect at several geographical locations to exchange messages.

network owners decide which standards to adopt. There is no overall authority to enforce standards or to govern interconnection business practices. Everything is negotiated between the network and internet owners. Who pays for the Internet? You do. Users pay ISPs, who work out arrangements with other ISPs to deliver your messages. You probably pay around $30 per month to your ISP. Businesses pay thousands or millions of dollars annually. With rare exceptions, no government money sustains the Internet.

Test Your Understanding

4. a) Is the Internet a single network? Explain. b) What is the role of ISPs? c) Who controls the Internet? d) Who funds the Internet?

Netflix Dives into the Amazon

You know personally how individuals use the Internet. The corporate experience is often very different. We will illustrate this by talking about how Netflix uses the Internet. Netflix is a commercial streaming video service with tens of millions of customers around the world. Streaming video places a heavy load on network capacity. For a two-hour high-definition movie, Netflix must deliver five million bits (1s or 0s) each second. This is a total of nine gigabytes for that one movie. On any given night, Netflix accounts for roughly a third of the Internet traffic going into U.S. homes.

Requirements Users expect high video quality, and they will not tolerate delay or unreliability. The Internet was not designed for these requirements. The Internet is a "best effort" delivery system that often has insufficient speed and reliability and that often has too much delay for Netflix users. Netflix had to overcome these limitations.

The Internet is a "best effort" delivery system.

Video streaming also requires vast amounts of server processing capacity beyond the demands of actual streaming. Each movie must be **transcoded** into many streaming formats, and when a customer requests a movie, streaming servers have to select the best transcoded format for that particular customer.

In addition, at the heart of Netflix's business plan is an application that creates personalized viewing suggestions for individual customers. This requires the analysis of extensive data about the customer's viewing habits and the choices of other customers with similar viewing profiles.

Outsourcing In 2008, when Netflix was only delivering movies through mailed DVDs, the company suffered a crippling server outage that stopped shipments for several days. That was a wake up call for Netflix. Management realized that reliability would be critical for the online delivery it would soon introduce. It also realized that while Internet delivery would become its core business, managing servers would not. Rather than developing the expertise needed for the complex server technologies the company needed, Netflix decided to outsource server operation to a company that could meet Netflix's high requirements for capacity, reliability, and agility in responding to sudden demand changes.

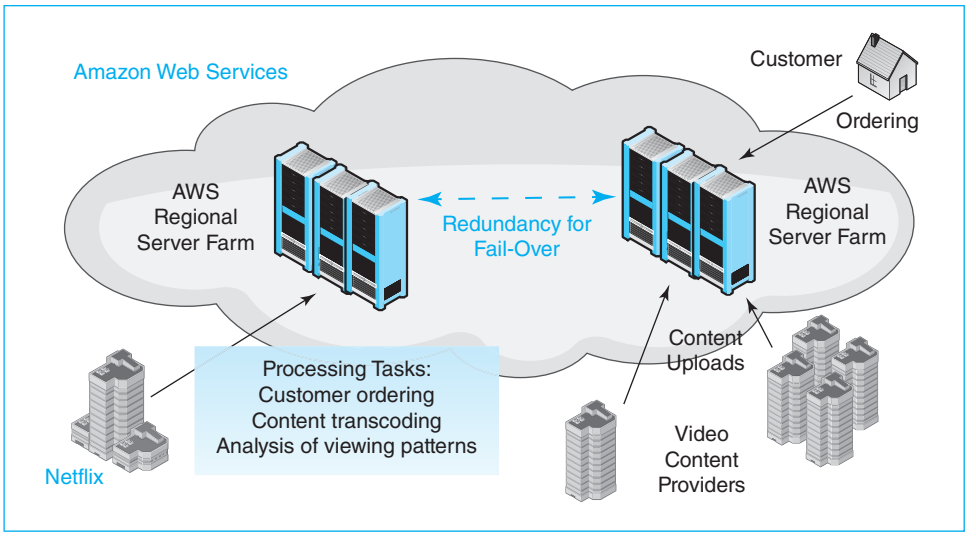

FIGURE 1-4 Netflix and Amazon Web Services (AWS)

Netflix chose **Amazon Web Services (AWS)**. Amazon had leveraged its expertise in managing vast server farms for its e-commerce needs into a cloud service that customers like Netflix could use without worrying about how the servers are operated. Figure 1-4 shows that AWS's enormous server farms had the capacity that Netflix needed for customer ordering, transcoding, and the analysis of viewing patterns. In addition, Amazon had multiple regional server farms with high fail-over capabilities. Even the loss of an entire server farm would not disrupt service for more than minutes. This brought the reliability that Netflix customers demanded. Netflix customers today log into an AWS server to order videos and to take care of other business transactions with Netflix. There are many login servers, and AWS automatically routes the user to one of them. Movie content providers upload their video directly to AWS. Netflix then transcodes the contents into many versions optimized for particular combinations of network speed and customer equipment.

Content Delivery Netflix uses AWS to store more than one petabyte of movie content in multiple locations. However, Netflix handles content delivery itself. Figure 1-5 shows how Netflix delivers video content to individual customers. Netflix calls this **content delivery network (CDN)** *Open Connect*.

To stream movies to users, Netflix created its own webserver appliances. Each is a relatively small box that can fit into a standard 19-inch (48-cm) wide equipment rack. The Open Connect appliance is seven inches (18 cm) high and two feet (61 cm) deep. Although small in size, it holds about 100 terabytes of data on 36 hard disk drives. The processor is fast enough to stream movies simultaneously to between 10,000 and 20,000 customers. Netflix updates these CDN servers about once a year with newer hardware to increase their capabilities.

Figure 1-5 shows that *Open Connect* is a network on the Internet. It can peer with (connect to) the ISP of a customer. The CDN boxes are placed at the peering point, so

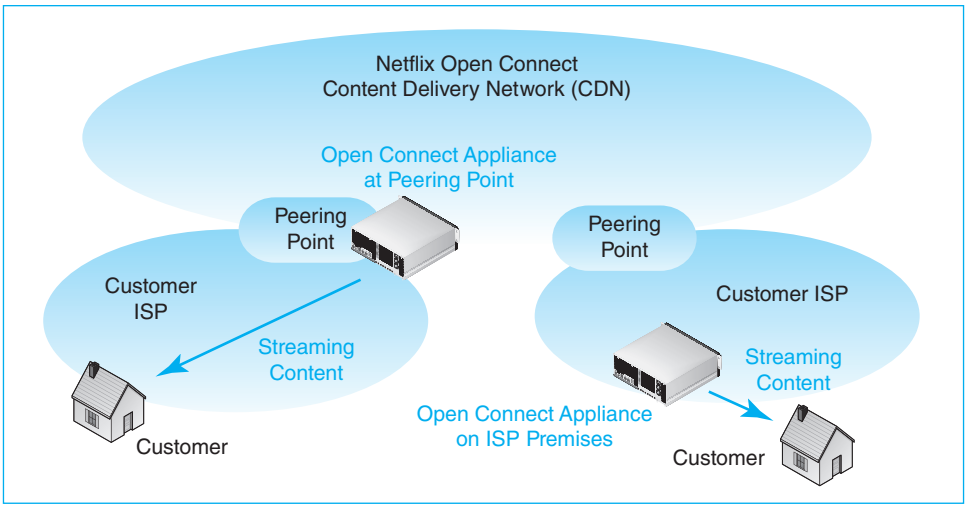

FIGURE 1-5 The Netflix Open Connect Content Delivery Network (CDN)

that traffic only travels the relatively small span of the customer's Internet service provider network. In many cases, ISPs reduce delays further by placing the appliance on their own premises, at a location near the final customer. ISPs tend to like this approach because it reduces traffic flowing across their network. The streaming traffic only goes the short distance from the nearest ISP physical location to the customer.

With only 100 TB of data storage, Open Connect appliances can only handle a portion of Netflix's 1 petabyte of content. Consequently, Netflix uses sophisticated analysis to identify the 100 TB of content most likely to be demanded by customers. It installs this content on the individual CDN servers. Of course, customer interests change rapidly, so this content has to be refreshed daily. During quiet periods in demand each day, Netflix deletes content declining in popularity and installs content of increasing demand.

Test Your Understanding

5. a) List Netflix's content delivery requirements. b) What is transcoding? c) Why does Netflix make many transcoded versions of each movie? d) How does Netflix use AWS? e) How do content delivery networks reduce streaming delays to customers?

Virtualization and Agility

Figure 1-6 shows that AWS uses virtualization to turn each physical server into several **virtual machines (VMs)**. Each VM is a software process running on the physical server. However, it acts like a real server in its connections with the outside world. It has its own IP address as well as its own data. It is even managed like a real server.

Using virtual machines gives an organization **agility**, which is the ability to make changes quickly—even very large changes. For example, Amazon can move VMs quickly from one physical server to another simply by transferring their files. It can even move VMs to servers quickly to different regions of the world. In addition, new VM **instances** (specific virtual machines) can be added in seconds. In fact, a company can **spawn (instantiate)**

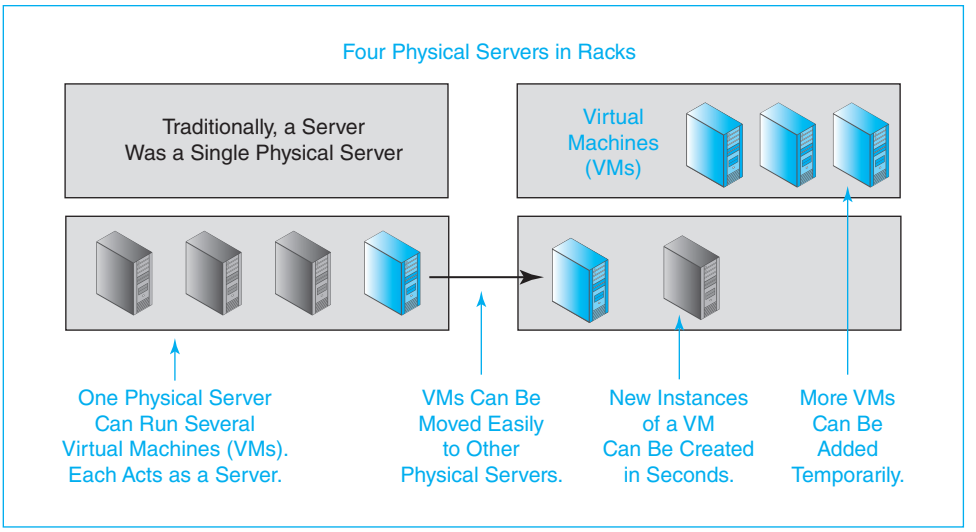

FIGURE 1-6 Server Virtualization through Virtual Machines

many copies of the same virtual machine at once, in no more time it takes to spawn a single VM instance. Physical servers offer nothing like this degree of agility. To make virtualization even more attractive to customers like Netflix, AWS provides a simple self-service application for customers to use to add new instances and do many other things themselves.

Content delivery is not the only way Netflix uses Amazon Web Services. Transcoding each movie into a hundred or more versions for delivery is an enormous task. Whenever Netflix needs to transcode a movie, it spins up (spawns) a large number of VMs, splits the work up among them, processes the data in parallel, and then spins them down. Providing customized viewing recommendations to subscribers also requires an enormous amount of processing power because it uses an extremely sophisticated analysis of individual user viewing practices and the viewing practices of people who have viewed similar movies. This recommendation system also requires Netflix to spin up large numbers of servers for short periods of time. Even in content delivery, the ability to spawn and kill VMs quickly is critical. During peak evening viewing time in the United States, Netflix spins up many additional VMs for content delivery. It spins them down later to save money.

Test Your Understanding

6. a) Distinguish between physical servers and virtual machines. b) What can be done with virtual machines that would be difficult to do with physical servers? c) What is VM instantiation? d) How does Netflix use the agility offered by Amazon Web Services?

Infrastructure as a Service (IaaS) and Software as a Service (SaaS)

Amazon is a **cloud service provider (CSP)**. Figure 1-7 illustrates this concept. We saw earlier that the Internet and other networks are depicted as clouds. The figure shows that CSPs also operate their services opaquely, forming a second layer of cloud.